CCS PERFORMANCE MODEL - INITIAL RESULTS

Preface

The following is a description of the initial performance model, results, and future plans. The initial model uses an analytical queueing theory approach and provides mean value performance predictions for Release 1 delay, throughput, and utilization metrics for steady-state conditions. This modeling methodology was chosen for the initial model because: 1) it can provide relatively accurate results and 2) model development and execution times are small fractions of the corresponding times for a discrete-event simulation model. This approach allows the primary performance problems of the architecture to be identified early enough for system architects and designers to make modifications in a timely manner. However, the methodology has limitations in terms of modeling some system complexity and providing analysis of peak conditions. For these reasons we also plan to develop a discrete-event simulation model.

The initial model reflects the functions, hardware configuration, and workloads for Release 1 plus some of those fore Release 2. The initial results indicate there are no performance problems. However, there are a few potential bottlenecks that should be watched, especially since all of Release 2 functionality is not included. They are the Firewalls, FEP CPUs, and the Application Server Disk.

This report is organized into descriptions of the CCS architecture modeled, hardware service rates, CCS application service demands, CCS workloads, and initial performance predictions. An attachment provides the original data collection templates and data from which model parameters were derived.

1. CCS Architecture

The initial model reflects the Release 1 configuration plus some of Relase 2. As shown in Figure 1, the star architecture consists a 100 Base-T Switch that interconnects:

- Front End Processor (FEP)
- Core Data Server (CORE DS)
- Backbone Data Server (BB DS)
- Application Server (APPL SVR)
- Core Firewall (CORE FW)
- External Web Server (EXT WEB)
- Internal Web Server (INT WEB)
- External Hub (EXT HUB)
- Internal Hub (INT HUB)

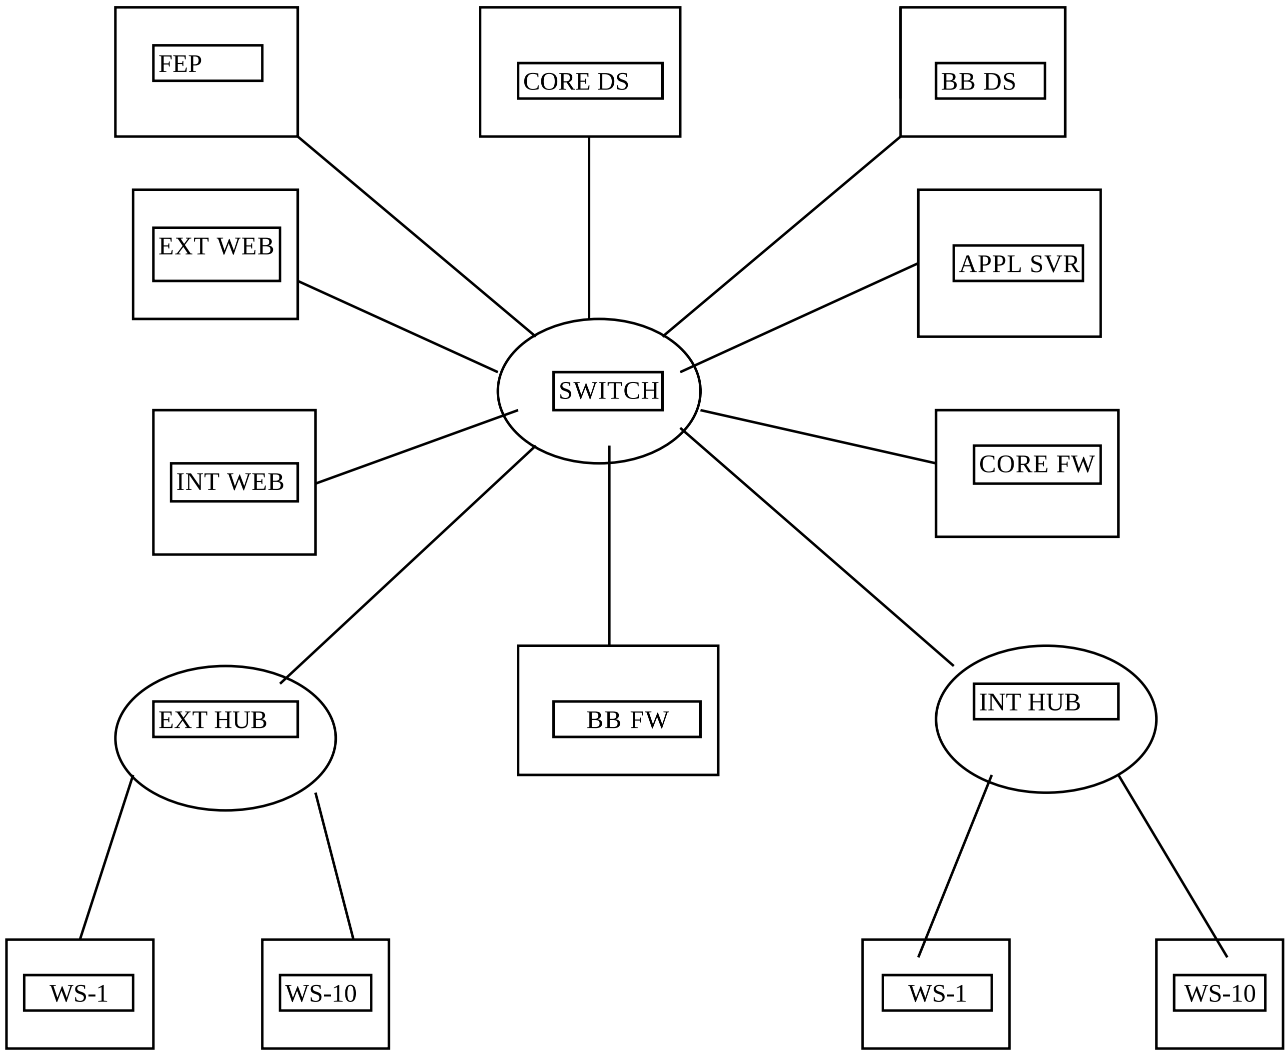

FIGURE 1 RELEASE 1 ARCHITECTURE

The data rate of the physical connections between the Switch and the nodes listed above is 100 Megabits per second (Mbps) except for the connection to the External Hub which is at ATM rates. A rate of 155 Mbps is assumed for that link. The two Hubs in the model are assumed to interface ten workstations each connected by 10 Base-T links with data rate of 10 Mbps.

Note that the Cisco Routers have not been included. They are not expected to contribute significantly to thread delays or to be performance bottlenecks. The external switch was also not included. The internal switch was included for modeling convenience; neither device is expected to have any impact on performance.

In addition to the current version of the architecture described in Table 1, the previous generation of the architecture was modeled to evaluate performance differences between the architectures. The main differences in the old and new generations are how the Web and Application Servers connect to other components and an additional Firewall for external users. In the older generation the Web and Application Servers were connected by 10 Base-T links to the Analysis LAN Hub; in the newer generation, they are connected by 100 Base-T links to the Switch. The results showed no real difference between the two architectures' performance.

2. Hardware Service Rates

The analytical model is defined in terms of a network of queues and servers that correspond to CPUs, disks, and network links in the CCS architecture. Each server is characterized by a service rate determined from vendor information. Service rate is combined with application service demand in the model to determine service times for workload transactions as they traverse the network of queues. Hardware service rate parameters used in the model are given in Table 1 for nodes and in Table 2 for links.

CPU service rates are characterized by SPECint92[1] rates published by the vendors. CPU data was gathered from Sun, SGI, and HP Web pages. The SPECint92 benchmark was chosen for several reasons: vendors have results for most all of their CPUs; it provides a common standard for CPU comparison; and the units are roughly comparable to MIPS (millions of instructions per second) for integer type code. The last reason is important if service demand estimates are provides in terms of number of instructions per software CI execution. The SPECint92 benchmark will not necessarily match the instruction mix for the CCS applications; however, it does provide a reasonable basis for estimating service times. If service demand estimates is provided in terms of measurement data from a CPU with a known SPECint92 rate, the corresponding service demand on another CPU is determined from the ratio of the two CPU rates. An overhead of 25% was added in the model to account for operating system and other system software that results in background contention for applications and other in-line thread software executing on the same CPU.

Disk service rate is characterized by three parameters: average seek time, average rotational latency time, and data transfer (i.e., read/write) rate. These three parameters will determine the effective service rate of the disk. Disk data was gathered from IBM, Seagate, and Quantum Web pages. These types of disk drives are used in the SGI and Sun computers in the CCS architecture. These parameters are combined

[1] SPEC, the Standard Performance Evaluation Corporation, defines standardized performance benchmarks that are used widely by computer vendors.

with disk service demand estimates to give disk service time per software CI execution. For a given amount of data to be read or written by a CI per execution the total disk service time is determined from:

(seek + rotation + (bytes per access/transfer rate)) * number of accesses

The average seek time value for cases where successive write operations are at most one track apart can be lower than the average reported by a vendor. In some cases for disk writes the seek time may be partially or completely masked by an SCSI bus data transfer to the disk controller. From a modeling perspective this is considered to be logically equivalent to a seek delay but is characterized by the bus rate instead of the disk seek rate. Rotational latency delay is determined from the spin rate of the disk. A typical value reported by vendors for an average is the time for the disk to spin one half way around. For RAID disks the model assumes that a group of *n* drives where data are "striped" can transfer data at the rate of *n* times the transfer rate of a single drive.

The amount of data per disk access has an impact on the throughput of a disk since it will determine the number of disk accesses for a given amount of data read or written by a process. The baseline assumption for the model is 1.0 MBytes per access.

Communication link service rate is simply the data transfer rate. The service time of a network transfer is determined from the amount of data per transfer, including protocol overhead for headers, divided by the transfer rate of the link. The model includes overhead for link layer protocols for Ethernet and ATM, as appropriate, and for TCP/IP for all of the links. Overhead for the ISP protocol is included in the application data size estimates. The overhead accounts for headers, etc. for user data packets as well as background packets for functions such as acknowledgment and flow control.

Table 1 Node Hardware Service Rate Parameters

Node	CPU	Disk
FEP	HP 712/100; 4 CPUs 117.2 SPECint92/CPU	1 disk drive Average seek = 8.5 msec Average rotation = 4 msec Transfer rate = 6.5 MB/sec
Core Data Server	SGI XL2; 2 CPUs 303 SPECint92/CPU	20 disk drive RAID Average seek = 8.5 msec Average rotation = 4 msec Transfer rate = 5*6.5 MB/sec
Backbone Data Server	SGI XL; 4 CPUs 303 SPECint92/CPU	20 disk drive RAID Average seek = 8.5 msec Average rotation = 4 msec Transfer rate = 5*6.5 MB/sec
Internal Web Server	SGI Indy-2; 1 CPU 140.2 SPECint92/CPU	1 disk drive Average seek = 8.5 msec Average rotation = 4 msec Transfer rate = 6.5 MB/sec
External Web Server	SGI Indy-2; 1 CPU 140.2 SPECint92/CPU	1 disk drive Average seek = 8.5 msec Average rotation = 4 msec Transfer rate = 6.5 MB/sec
User Workstation	Sun Sparc 20/71 (typical) 1 CPU 125.8 SPECint92/CPU	1 disk drive Average seek = 8.5 msec Average rotation = 4 msec Transfer rate = 6.5 MB/sec
Core Firewall	Sun Sparc 20/71; 1 CPU 125.8 SPECint92/CPU	not relevant
Backbone Firewall	Sun Sparc 20/71; 1 CPU 125.8 SPECint92/CPU	not relevant
Application Server	SGI Indy-2; 1 CPU 140.2 SPECint92/CPU	1 disk drive Average seek = 8.5 msec Average rotation = 4 msec Transfer rate = 6.5 MB/sec

Table 2 Communications Link Service Rates

Node Pair	Transfer Rate (Mbps)
FEP - Switch	100
Core Data Server- Switch	100
Backbone Data Server- Switch	100
Internal Web Server- Switch	100
External Web Server- Switch	100
Internal User Workstation- Switch	100
Core Firewall- Switch	100
Backbone Firewall- Switch	100
Internal User Workstation - Internal Hub	10
External User Workstation - External Hub	10
Backbone Firewall - External Hub	155

3. CCS Application Service Demand Estimates

CCS application service demand estimates and information on COTS software were collected from the Release 1 teams using the templates in Attachment 1. The estimates were used with the hardware service rate information to derive transaction service times for the CCS workloads in the model. The accuracy of any performance model, whether simulation or analytical, is very dependent on the accuracy of service demand and arrival rate information. The level of confidence in the current service demand information is medium. Some of the estimates are based on measurement data and others are educated guesses. Table 3 gives the derived application service demand parameters used in the model. Processing service demand is given iin units of thousands of instructions per CI execution; Disk service demand is given in units of bytes of data read or written per CI execution; and Network service demand is given in units of bytes of data transmitted per CI execution. Note that the service demand parameters are based on the units of the transactions for the workloads. These are described in the next section.

Table 3 Service Demand Input Parameters

NODE	S/W CI	K INSTRUCTIONS	DISK IO (BYTES)	NETWORK IO (BYTES)
FEP	FE COTS ETR	8766986		
	Minor Frame Archive ETR	14611644	416880000	208440000
	FE COTS RT	42068.412		
	Minor Frame Archive RT	70114.02		70000
	Route Engr Data ISP	14020		10865
	Route Engr Data ISIS - RT CDF	21030		162500
	Route Engr Data ISIS - Recorded	21030		5200000
	Update Time	21.03		
FW	Firewall COTS	129.1248207		
CORE DS	File Server	1000	208440000	
BB_DS	ISP Server	28040		10865
	UNIX ISP Client	785.12		
	Record Stream Data - RT	2250	162500	
	Record Stream Data - ETR	2250	5200000	
	ORACLE - record data	750	1200	
	ORACLE - query data	750	401072	
	Merge TLM Files	2250	1600000	
	Archive Stream Data	150	54166.66667	
	Manage Data Requests	150		
	Perform Data Retrieval	3250	1000000	1000000
WEB SERVER	ISP Server	11216		1086.5
	UNIX ISP Client	785.12		
	JAVA ISP Client	785.12		
	WebServer- Data Analysis	1402	300000	300K applets/5 MB
	Web Server - RT display	1402	200000	200000
	GUI Server- Data Analysis	1402		5000-User/200-Appl_svr
	GUI Server- display setup	1402	10000	10000
	GUI Server- RT display	1402		10865
APPL SERVER	Analysis Job Management	1200		
	Build Analysis Request	1200	1000000	
	Generate Analysis Product	3880		
	Manage Analysis Data	1240		200
	Convert & Deliver Product	1200	20000000	10000000
	PVWAVE	1720	13000000	
USER WS	Analysis Web Browser	701000	300000	400
	RTD Web Browser - setup	1402000	200000	
	RTD Web Browser - disply	1402000		

Many of the CPU service demand estimates is based on measurements of Release 1 prototype or legacy software. For all of these cases a percentage of CPU utilization was reported. To determine transaction service time from the utilization, either the transaction arrival rate or the period of time during which the utilization was measured must be known. In some cases this information was available and in other cases educated guesses were made. All service demand estimates were standardized to Instructions per Software CI execution.

For the cases where prototype Release 1 prototype software was measured, the COTS infrastructure overhead is included. To the extent that the prototype includes dynamic object creation and destruction, the overhead associated with the object-oriented development methodology will be included in the measurements. For other cases the additional service demand for the overhead functions has been estimated based on past experience.

Service demand for the Checkpoint Firewall software was estimated from independent benchmark results published by *Data Communications* magazine and is given on a per IP packet basis.

Estimates for the amount of disk I/O per Software CI execution are based on designers' knowledge of how much data different types of transactions involve or on data rates and quantities from the legacy system. For the Ingest and Archive functions data quantities are based on telemetry rates and should provide accurate estimates. For analysis request transactions the quantities are based on knowledge of the legacy system and on knowledge of the infrastructure software. The range of the amount of data for an analysis response is very large. The values used in the model were reported as typical values. Different values need to be examined.

Estimates for the amount of data per network transfer have a basis similar to the disk estimates. In many cases the amount of data retrieved from a disk is the amount of data transferred over a network link.

4. CCS Workloads

The CCS workloads in the model were derived from the Top Down Analysis (TDA) system threads. These threads reflect the functions and physical architecture for Release 1. Several functions and thread segments have since been moved to Release 2. These functions and thread segments have been retained in the model where service demand estimates were already available. Thus, the model actually represents a Release 1+ CCS. The model workloads are summarized in Table 4 and briefly described below.

RT MF Archive - Real Time Minor Frame Archive workload represents the NASCOM blocks that arrive to CCS over a 56 Kbps link with fill data. These blocks, with the fill data, are processed by the FEP and sent to the Core Data Server. The model assumes that the blocks are buffered in FEP RAM and then sent to the Core Data Server when the buffer fills. The initial assumption for buffer size is 1 MB.

ETR MF Archive - Engineering Tape Recorder Minor Frame Archive workload represents the NASCOM blocks that arrive to CCS over a 1.544 Mbps link with fill data. These blocks, with the fill data, are processed by the FEP to forward-order the data and are written to a local disk before being sent to the Minor Frame Archive on the Core Data Server. Note that the Solid State Recorder (SSR) MF Archive thread does not need to be forward-ordered and could operate like the RT MF Archive thread. This workload is not represented in the current model because a thread definition is not available. It can be added later.

RT Archive - Real Time Archive workload represents the archive of real time CDF engineering data on the Backbone Data Server. The model assumes that CDF packets are packaged in groups of 40 and sent from the FEP approximately once per second. The initial assumption on the size of the CDF container is 1.3 Mbits since at the frequency of one per second this matches the 1.3 Mbps rate that has been documented.

Recorded Archive - Recorded Archive workload represents the archive of recorded (ETR or SSR) CDF engineering data on the Backbone Data Server. Because the recorded data stream has an data rate of 1024 Kbps, as compared to the real time stream rate of 32 Kbps, the model assumes that the size of the CDF container sent from the FEP to the Backbone Data Server is 1024/32 = 32 time larger than the real time CDF container, or 41.6 Mbits.The frequency of the workload is three times per day.

RT ISP Display - Real Time ISP Display workloads for internal (i.e., local to Building 23 Analysis LAN) and external CCS Users represent the stream of once per second updates from the FEP to User Workstations via the Backbone Data Server. The initial assumption for the amount of data is for 300 mnemonics per update. In the ISP format each mnemonic is represented by 36 bytes, and each ISP packet has 65 bytes of header overhead. Thus, the amount of data per ISP packet is 10, 865 bytes. Each of the ten internal and ten external CCS users in the model is assumed to receive 1/10, or 30 mnemonics, per update.

Setup RT ISP Display - Workloads are included for internal and external CCS users setting up the real time ISP displays with the Web Server. The initial assumption is each user performs two setups per day.

Analysis Request - Workloads are included for internal and external CCS users making Analysis Requests at the rate of 20 requests per day per user. The model assumes 10 internal users and 10 external users initially. The flow of these transactions as shown in the table below was abbreviated because of the tabular format. The flow involves interaction between the user and a Web Server initially to download the JAVA applets that are used in making the requests. The service demands associated with this interaction are included in the workload flow in the model. The tabular flow also does not show the return path from the Backbone Data Server through the Application Server and the Web Server to the User Workstation. This part of the flow is also represented in the model workload. The model initially assumes that the Backbone Data Server returns 1 MB of data to the Application Server which produces a 10 MB analysis product that is sent to the Web Server, and 5 MB are sent to the User Workstation. The data amounts used for the response have been described as a typical analysis product. The model assumes that these quantities are averages.

Time - A workload is included for the universal time input to the FEP and the Core Data Server. This workload could be extended later to represent time distribution throughout the system. For both FEP and Core Data Server computers the model assumes that the time data is maintained on a separate disk drive from engineering data.

Table 4 CCS Workloads

Workload	Frequency	Input Data Quantity	Execution Flow
RT MF Archive	20 hrs/day	NASCOM 4800 bits/block @ 56 Kbps	NASCOM -> FEP -> Core DS
ETR MF Archive	3/day	NASCOM 4800 bits/block @ 1.544 Mbps for 18 min	NASCOM -> FEP -> Core DS
RT Archive	1/sec	CDF container = 40 IP packets = 1.3 Mb	FEP -> BB DS
Recorded Archive	3/day	CDF container = 41.6 Mb	FEP -> BB DS
RT ISP Display - Internal	1/sec	ISP packet = 10,865 Bytes	FEP -> BB DS -> Internal Web Svr -> Internal WS
RT ISP Display - External	1/sec	ISP packet = 10,865 Bytes	FEP -> BB DS -> External Web Svr -> External WS
Setup RT ISP Display - Internal	2/day/user	200 Bytes	Internal WS -> Internal Web Svr
Setup RT ISP Display -External	2/day/user	200 Bytes	External WS -> External Web Svr
Analysis Request - Internal	200/day	200 Bytes	Internal WS -> Internal Web Svr -> Appl Svr -> BB DS*
Analysis Request - External	200/day	200 Bytes	External WS -> External Web Svr -> Appl Svr -> BB DS*
Time Input	7/sec	32 bits	FEP & Core DS

* Analysis Request execution flow description is abbreviated in the table. See the text description above for the full path description.

5. Initial Performance Predictions

Analytical models have the limitation of being able to produce only mean value statistics for steady-state conditions. All of the initial results given below are of this type. Model prediction results for the initial set of input values collected from vendors, the CCS designers, and other sources are reported in Section 5.1, Baseline Predictions. Various parameters were varied to examine their performance impact. Sensitivity Analysis prediction results are reported in Section 5.2.

5.1 Baseline Predictions

Initial baseline predictions are summarized in Table 5 and Table 6 for node resource utilization and workload delays, respectively. Utilization for the network links is three percent or lower. Many links are utilized less than one percent. The potential bottlenecks in the architecture are the Firewalls, the FEP CPUs, and the Application Server Disk. The utilization and component delays for these resources are not problems for Release 1 functionality and loads; however, they would be the first resources to become saturated in the architecture.

The Core Firewall average utilization is 12.8 percent. The peak situation will occur during when recorder dumps are being archived on the Backbone Data Server. The maximum throughput of the Checkpoint Firewall, according to our preliminary estimate, is around 11 Mbps. This scenario was examined and is discussed in Section 5.2 below. Note that the potential Firewall bottleneck can be alleviated by load-balancing the workloads over multiple Firewalls.

FEP processor utilization is at 13.6 percent per CPU for a four CPU configuration. The model does not reflect a playback (from the Minor Frame Archive) workload at this point. If this workload can occur simultaneously with recorder dumps, then there is potential for a peak utilization problem.

For the ETR Minor Frame Archive workload over 93 percent of the delay is attributable to the FEP CPU and disk. Core Data Server delay contributes only two percent of the delay; the remainder of the delay is due to the network.

Delay for setting up a real time display is due almost exclusively to the workstation software with only one percent of the delay due to the Web Server and network.

For the Analysis Request workload 56 percent of the delay is attributable to the User Workstation; 20 percent is attributable to the networks; and 22 percent is due to the Application Server. The delays for the Web Server and the Backbone Data Server are two percent of the total delay. The service demand and arrival rate parameters for this workload were varied in several sensitivity analyses as described below.

Table 5 Baseline Utilization Predictions

Node	CPU Utilization	Disk Utilization
FEP	13.6% per CPU	9%
Core Data Server	< 1% per CPU	< 1%
Backbone Data Server	6.1% per CPU	3.1%
Internal Web Server	5%	< 1%
External Web Server	5%	< 1%
User Workstation	6.4%	< 1%
Core Firewall	12.8%	N/A
Backbone Firewall	8%	N/A
Application Server	< 1%	7.8%

Table 6 Baseline Workload Delay Predictions

Workload	Delay
RT MF Archive	1.5 seconds/1 MB
ETR MF Archive	6.7 minutes per recorder dump
RT Archive	< 1 second/CDF Container
Recorded Archive	1.1 seconds/CDF Container
RT ISP Display - Internal	1 second/update
RT ISP Display - External	1 second/update
Setup RT ISP Display - Internal	12 seconds/setup
Setup RT ISP Display - External	12 seconds/setup
Analysis Request - Internal	29.5 seconds per request
Analysis Request - External	29.5 seconds per request
Time Input	<1 second/input

5.2 Sensitivity Analyses

Several sensitivity analyses were performed to examine the performance impact of changing parameter values. The follow cases were examined:

- Software CI service demand
- Average disk access time
- Amount of data per disk access
- Number of striped disks
- Throughput of recorder dumps
- Analysis Request Arrival Rate
- Analysis Product Size

Software CI Sensitivity

This sensitivity analysis was performed because of the uncertainty of some of the estimates for software service demand. The values derived from the design team inputs were varied to six times the baseline values to examine the impact to performance. From the baseline CPU utilization results (see Table 5 above) it can be seen that for a multiplier greater than seven, the FEP CPU will have utilization approaching 100 percent. All other parameter values in the model were not changed for this sensitivity analysis.

Figures 2, 3, and 4 show workload delay for the most-affected workloads. For Analysis Request and Real Time Display Setup workloads, shown in Figure 2, a large majority of the delay is in the user workstation. The growth in delay for these workloads is linear. Because the workstations are single user devices, there is no queueing to cause the curves to depart from linearity. The Real Time Minor Frame Archive workload delay shows the affect of queueing at the FEP CPUs.

The growth in the workload delays shown in Figure 3 is due to service and queueing times at the FEP CPUs. At a service demand multiplier of two, Real Time Display updates take over two seconds which is not fast enough to keep up with the one second update rate.

The delay for the Recorded Minor Frame Archive shown in Figure 4 is very sensitive to the FEP software service demand. The FEP software CIs for this workload are the Veda COTS software and the Minor Frame Archive CI. This latter CI could be a candidate for moving to the Core Data Server since its CPU has very low utilization.

As shown in Figure 5 CPU utilization increased by the same factor applied to software service demand. Utilization is a linear function of arrival rate and service demand. Multiplying all of the service demands of a resource by a factor of *n* will results in a utilization of *n* times the baseline utilization. Queueing delay is not a linear function of service demand, and workload delays will not necessarily increase linearly.

While there is no basis at this point to believe that the Software CI service demand estimates are off by a factor of two or more, these results point to the need to validate those estimates.

Figure 2: Software Service Demand Sensitivity - Dela

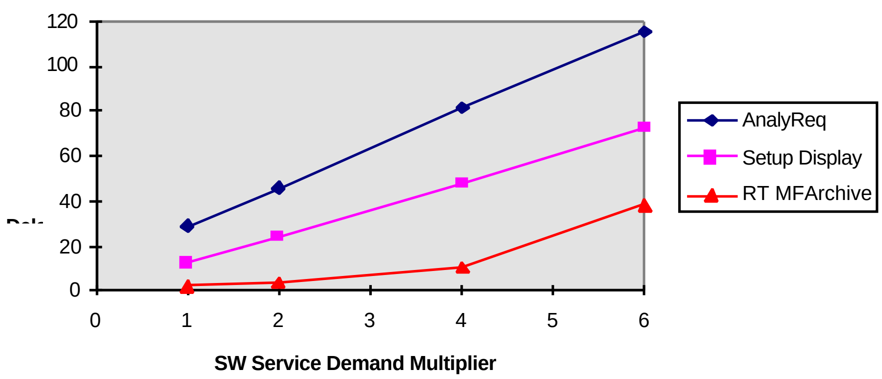

Figure 3 Software Service Demand Sensitivity - Dela

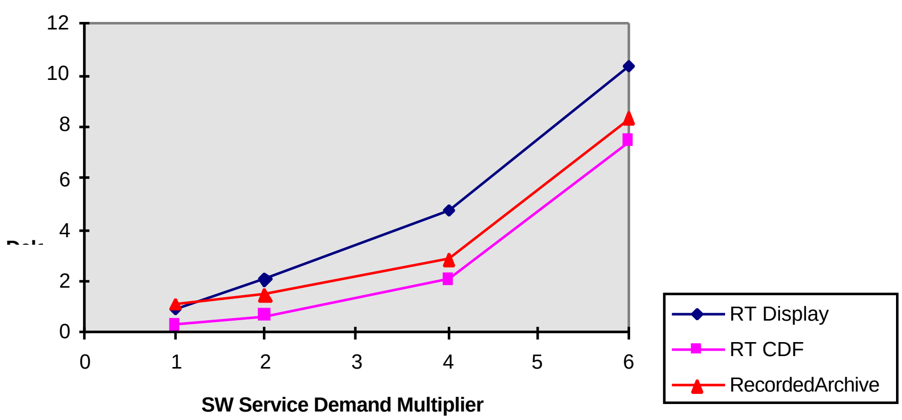

Figure 4: Software Service Demand Sensitivity - Dela

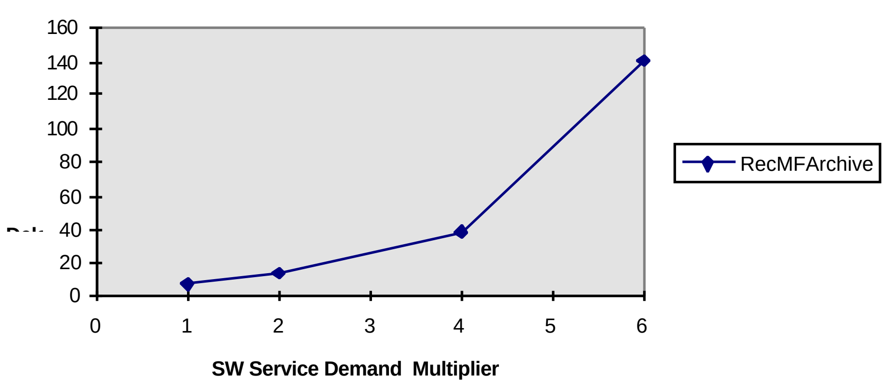

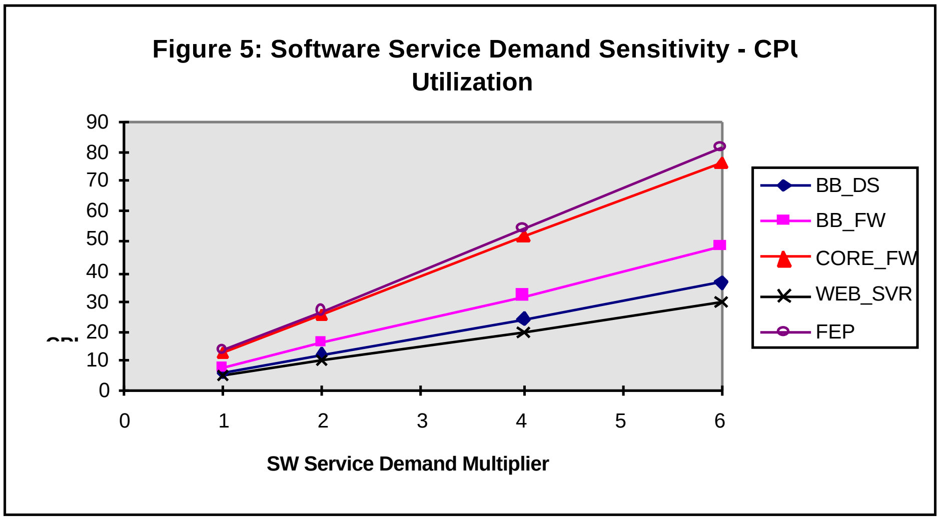

Figure 5: Software Service Demand Sensitivity - CPU Utilization
90
80
70
60
50
40
30
20
10
0
0
1
2
3
4
5
6
SW Service Demand Multiplier
BB_DS
BB_FW
CORE_FW
WEB_SVR
FEP

Average Disk Access Time Sensitivity

A value of 12.5 milliseconds for disk access time was used for the baseline case that produced the results reported in Section 5.1. Additional cases were examined for 4 and 8 milliseconds. All cases assume that the average rotational latency of the disk is 4 milliseconds which is inclusive in the access time Workload delay and disk utilization are given in Figures 6 and 7, respectively. This parameter can affect delay and utilization significantly; however, as can be seen by the results the disk utilizations are so low that there is no perceptible differences in delays. The Application Server disk shows the most potential for being sensitive to access time. Disk access time is the main part of disk delay for Analysis Requests.

Release 2 functions and loads could change the outcome of this sensitivity analysis. These will be added, and the disk access time parameter will be refined based on measurement data when the SGI Challenge servers at the co-location facility are configured with the RAID disks.

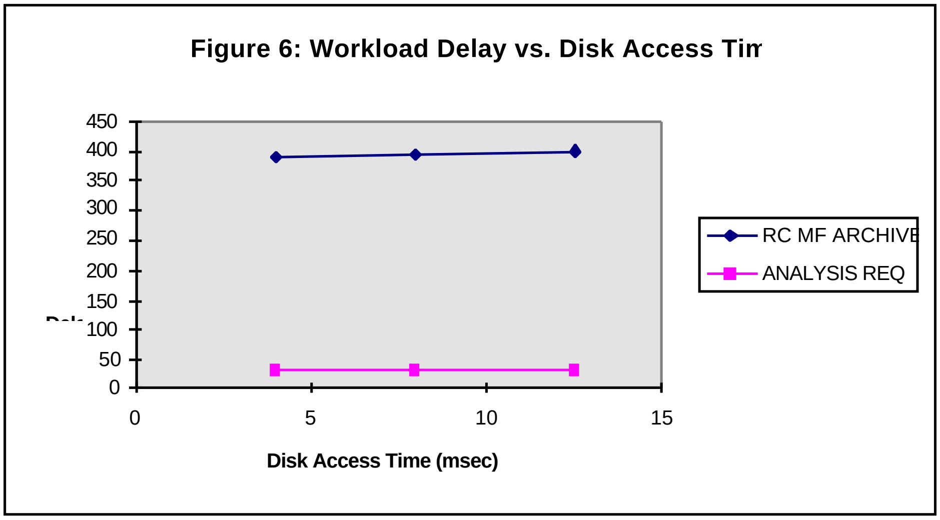

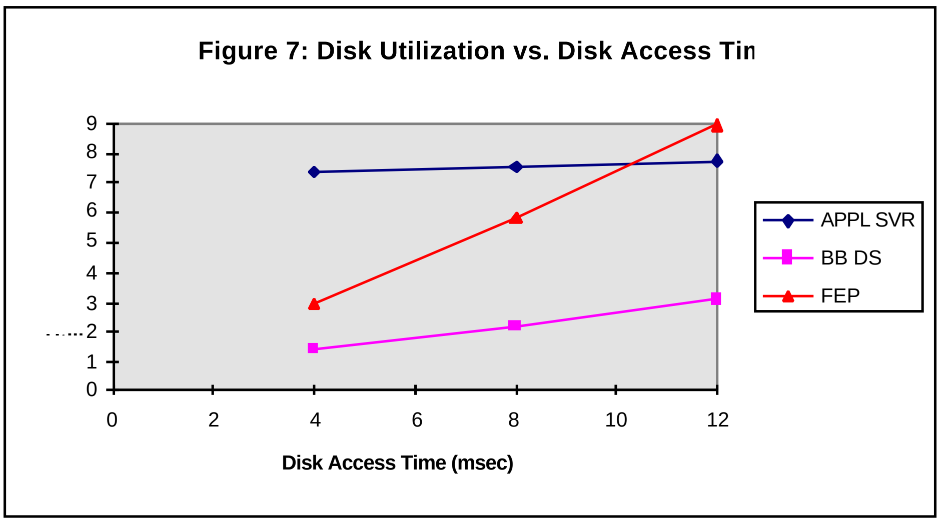

Amount of Data Per Disk Access Sensitivity

This parameter reflects how much data will be buffered before it gets written to disk and how much data is read from a disk for each access. This level of hardware design has not been decided at this time. Like disk access time, the amount of data per disk read or write is an important performance parameter. Disk access time (i.e., seek + rotational latency) is typically more significant than the read or write times. Therefore, by limiting the number of accesses for a given amount of data, disk delay and utilization can be reduced. Larger amounts of data per access translates into fewer accesses and higher disk data throughput. The baseline case assumption is 1 MB per access. A more pessimistic case of 0.5 MB per access was examined. The workloads that are most sensitive to delay are Analysis Requests and Engineering Tape Recorder (ETR) Minor Frame Archive. At the loads for the Release 1 functionality workload delays increased only by three percent. The only disk that showed much of an increase in utilization was the Application Server disk, which had a relative increase of about eight percent. The addition of Release 2 functions and workloads could change these results.

Another sensitivity case examined writing the Recorded NASCOM blocks individually to the Minor Frame Archive on the Core Data Server disk. In this scenario the blocks were sent as they were processed by the FEP. This resulted in the disk being saturated (i.e., it could not keep up with the rate). This scenario is not likely for ETR data since it must be forwarded-ordered at the FEP which involves writing it to disk. Solid State Recorder (SSR) NASCOM blocks could possibly be written to the Core Data Server disk one at a time, but these results show it's not a good approach.

Number of Striped Disks

The baseline assumption for the number of disk drives striped in the Core Data Server and the Backbone Data Server is five drives. Design decisions for these disk configurations have not been made at this point. This parameter was examined for values of 10 and 15 striped drives. The results are shown in Figures 8 and 9, respectively, for Recorded Minor Frame Archive delay and Core Data Server and Backbone Data Server disk utilization. The delay and utilization decreases are slight at the Release 1 loads. The addition of Release 2 functions and workloads could change these results.

Figure 8: Disk Striping Sensitivity - MF Archive Dela

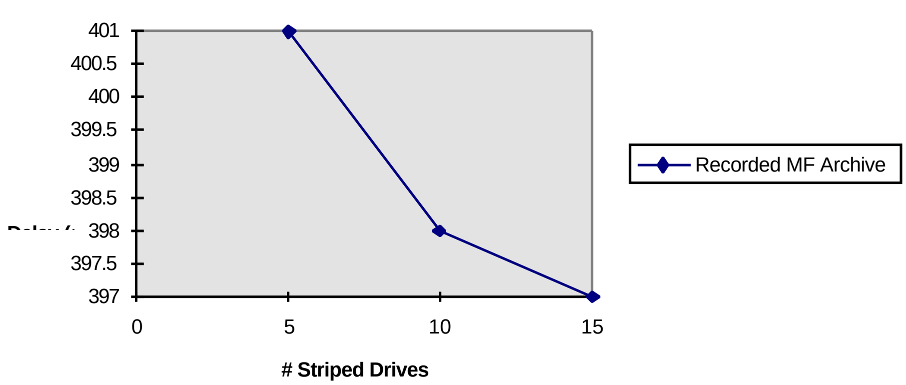

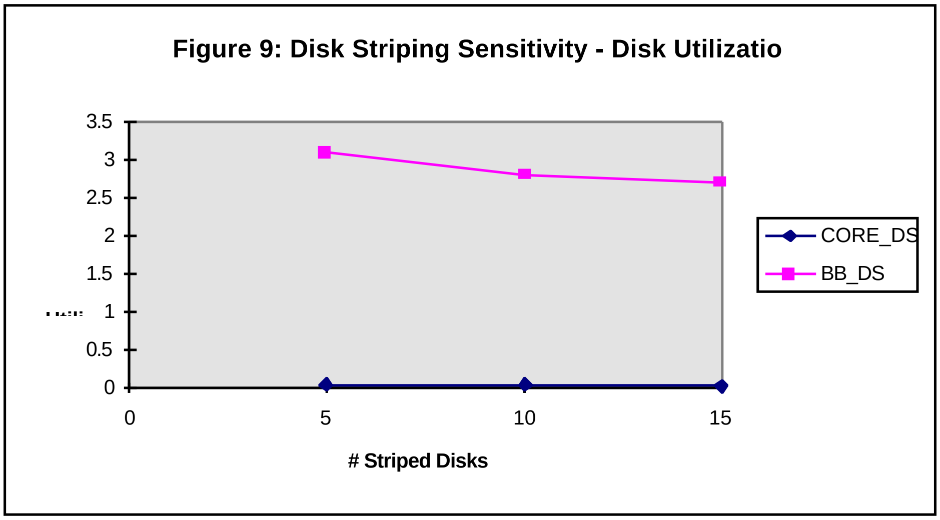

Recorder Dump Throughput

One system characteristic that has not been determined yet is the rate at which the recorder dumps can be sent from the FEP to the Backbone Data Server Archive. In this sensitivity analysis the Recorded Archive workload throughput rate was determined by the model instead of being an input. The assumption of the model representation is that the recorded data enters the system as fast as it can be handled (i.e., processed by the FEP, sent to the Backbone Data Server, and written to the Archive). This situation resulted in saturating the Core Firewall. The workload was throttled, by introducing an artificial delay, until the Firewall CPU was no longer saturated. The maximum throughput of the Checkpoint Firewall is around 11 Mbps total according to our initial estimates. The resulting maximum throughput for the Recorder Dump workload is around 9.5 Mbps. It is important to view this result as preliminary since the Backbone Data Server will have much more functionality and additional loads for Release 2. We will validate the performance capability of the Firewall by obtaining further performance data from the Checkpoint vendor or by benchmarking and will add Release 2 functionality and loads to the model in the future.

Analysis Request Arrival Rate

The baseline estimate for Analysis Request arrival rate assumes a nominal system state. This arrival rate will change as a function of the high-level states of the system (e.g., system anomalies, Servicing Mission, etc.). Loading scenarios will be developed for different system states for subsequent performance analysis. For this sensitivity analysis arrival rates of internal and external users were varied from 400 per day to 3200 per day to examine system performance impact. The size of the Analysis Product was held constant at the baseline value. The higher values of arrival rate range might correspond to HST system states such as spacecraft anomalies or Servicing Mission.

The results of the sensitivity analysis are shown in Figures 10 and 11, respectively, for mean response time and utilization for the Application Server disk, which is the most affected resource on this thread. Note that response time for both internal and external users differs by a tenth of a second at most for this range of transaction rates.

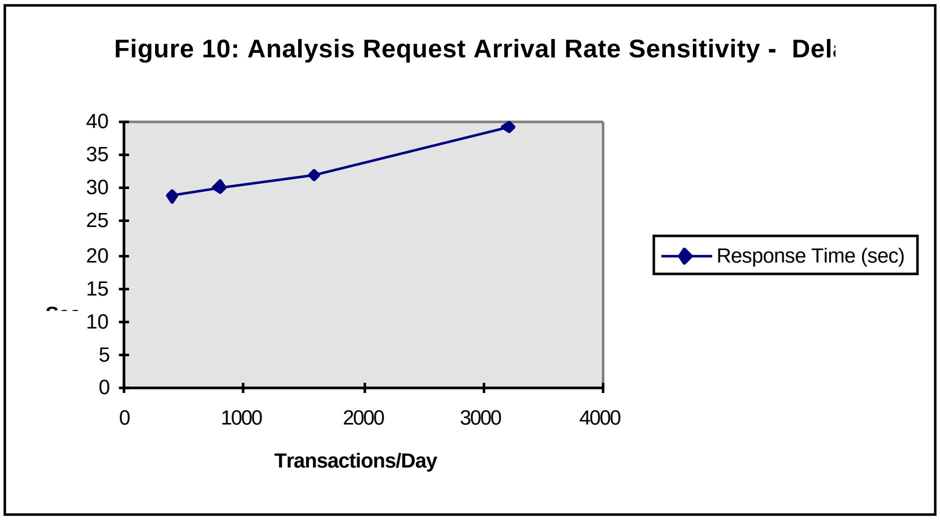

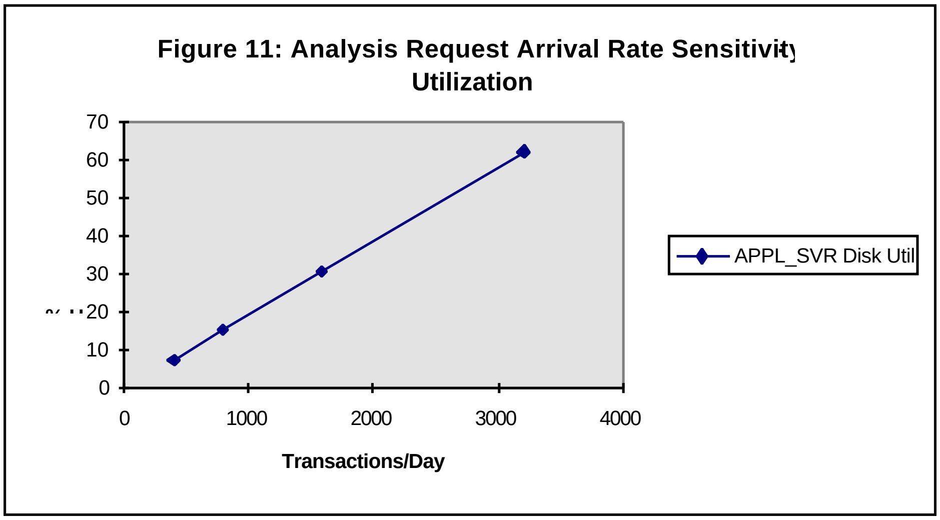

Analysis Product Size

This sensitivity analysis was motivated by the uncertainty of the estimates of the amount of data associated with a response to an analysis request, and the fact that the amount of data in a response is expected to have a large variance. The variance of the amount of data is related to high-level system states similar to Analysis Request arrival rate. This sensitivity analysis uses the baseline arrival rate for Analysis Request transactions but varies the amounts of data involved in constructing an Analysis Product. The assumption for Analysis Request transactions is that the Application Server produces an Analysis Product ten time larger than the amount of data it gets from the Backbone Data Server. Half the amount of data in the analysis product is sent by the Web Server to the user workstation. Analysis Product size was varied from 10 MB to 80 MB.

The results of the sensitivity analysis are shown in Figures 12 and 13, respectively, for mean response time and utilization for the Application Server disk and the Backbone Firewall CPU. These two resources are the most sensitive to the amount of data in a response. The most significant contributors to response time are the Firewalls, the user workstation, and the Application Server disk. The user workstation delay is constant throughout the range of Analysis Product size as expected. The communications delay increases but not as dramatically as the Application Server disk delay over the range. The disk delay in the Application Server at eight time the Analysis Product baseline size is eight times the baseline disk delay. The disk delay is due to five separate sets of disk accesses:

1. Analysis data from the Backbone Data Server is written to disk
2. PV WAVE Scripts are read from disk
3. PV WAVE software writes Analysis Products to disk
4. Convert & Deliver CI reads Analysis Products from disk
5. Convert & Deliver CI writes Final Analysis Product to disk

It may be possible to mitigate some of these delays by managing the analysis data and products in a buffer instead of having the in-line (to the thread) disk accesses (i.e., permanent retention of products could be performed after the response has been sent to the Web Server). This approach depends on the amount of data involved in the various stages of a response.

Figure 12 also shows response time for external users is more sensitive to Analysis Product size than for internal users. This is due to the Backbone Firewall device on the thread for external users.

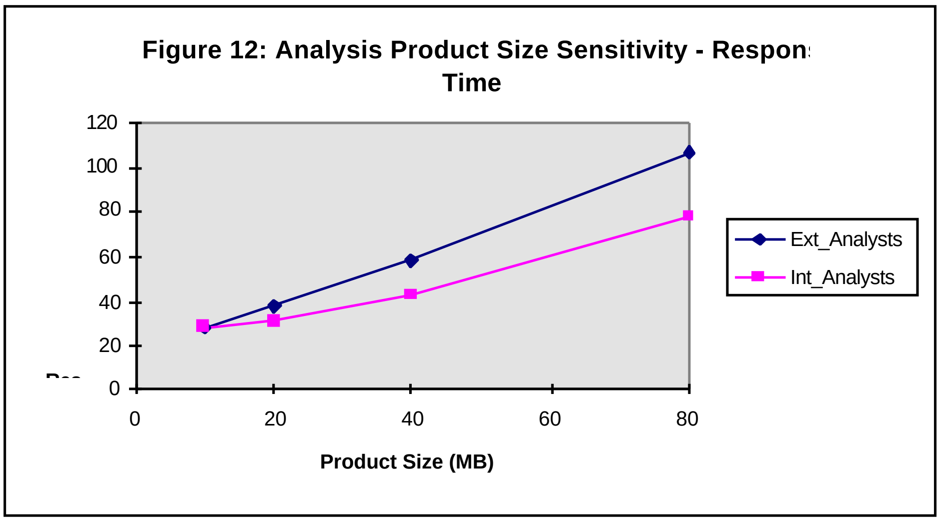

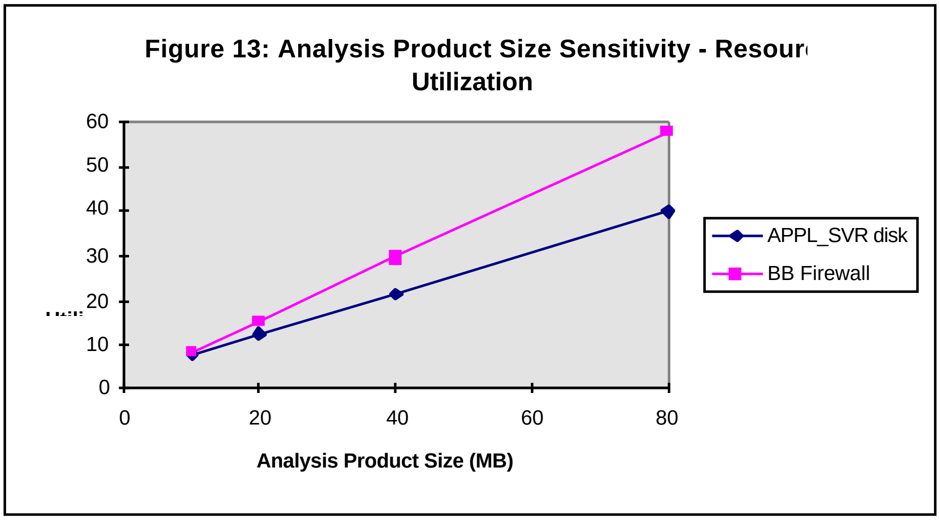
Figure 13: Analysis Product Size Sensitivity - Resour
Utilization
60
50
40
30
20
10
0
0
20
40
60
80
Analysis Product Size (MB)
APPL_SVR disk
BB Firewall

6. Future Plans

Future plans for the performance modeling activity include additional data collection, analytical model modification, simulation model development, and further performance analysis.

6.1 Data Collection

As indicated above, the accuracy of some of the model parameter values needs to be examined. Some of the parameter estimates will be refined as measurement data become available. We will depend on refined estimates from design team members for other parameters.

In addition to the data collected to-date, data need to be collected to characterize different system loading scenarios, the Release 2 threads and functions, architectural modifications, and the remaining system overhead functions like security auditing, enterprise management, and network management.

6.2 Analytical Model

The analytical model will be refined to incorporate the updated data collected and will be used to examine sensitivity cases such as:

- Analysts request rate and product size for different system states
- Architecture modifications
- Release 2 threads and functions
- Additional system overhead functions

6.3 Discrete Event Simulation Model

Finally, we plan to develop a discrete-event simulation model of CCS. This model will be higher fidelity and will provide distribution statistics of performance metrics as well as mean value, steady-state results that the analytical model produces. The discrete-event model could also be enhanced to examine equipment failure scenarios. Because of the time to develop and execute the simulation model, it is likely that only the final architecture will be modeled. Results from the simulation model will provide the definitive performance evaluation of CCS prior to testing but will also be used to validate the analytical models. The analytical models will continued to be used to examine sensitivity cases where timely feedback to the design teams cannot be provided using the simulation model.

Attachment 1 - Data Collection Templates

TEMPLATE 1: WORKLOAD FREQUENCY CHARACTERIZATION

Thread Id	Thread Name	Source	Frequency	Input Data Quantity
1	Time	GPS	7/second	32 bits
7	Data Analysis Request	CCSU	200/day TLM & File Catalogs 200/day PRD Catalog 50/day for RedBrick	200 bytes/request 200 bytes/request 1 KB/request
17	RT Data	NASCOM	20 hours/day	Stream: 32 kbps x 80-85 minutes
18	Recorded Data	NASCOM	3/24 hours	1024 kbps x 18 minutes
17a	RT Data Display	CCSU	1 second updates	Stream: 32 kbps x 86 minutes

TEMPLATE 2: SOFTWARE RESOURCE REQUIREMENTS CHARACTERIZATION

(page 1 of 3)

Node	**Software CI**	**Processing Demand**	**Memory**	**Disk I/ODemand**	**NetworkDemand**
FEP	Comm I/F	H/W	H/W	H/W	@ 56 kbps
	FE COTS	6 % @ HP7000	1 MB		@ 32 bps
	Minor Frame Archive	20 % @ SGI R4400 CPU (sum)	4 MB		
	Playback	10 % @ SGI R4400 CPU	2 MB		
	Achive	10 % @ SGI R4400 CPU	2 MB	1.6 mbps - Core DS disk	
	Route Engr Data	25 % @ SGI R4400 CPU (sum)	11.8 MB		
	ISP	10 % @ SGI R4400 CPU	6.8 MB		84/3.9 & 30/1.38[2]
	ISIS - CDP Send	15 % @ SGI R4400 CPU	5 MB		59/2.72 & 105/4.83
	Update Time	5 % @ SGI R4400 CPU	1 MB		
FEP CONFIG WS	Manage FE	10 % @ SGI R4400 CPU	4 MB	N/A	N/A
	Quicklook 6000 mnemonics	5 % @ SGI R4400 CPU	1 MB		300/13.8
	VEDA COTS	1 % @ SGI R4400 CPU	2.3 MB		
CORE DS	Time Server				
FIREWALL	Firewall	820 trans/min[3]			TCP/IP packets

[2] Packet rate/kbps for IP received packets & transmitted packets.
[3] Transaction size varied (1 KB, 256 KB); average 85 KB.

TEMPLATE 2: SOFTWARE RESOURCE REQUIREMENTS CHARACTERIZATION

(page 2 of 3)

Node	**Software CI**	**Processing Demand**	**Memory**	**Disk I/ODemand**	**NetworkDemand**
BACKBONE DS	ISP Server[4]	1.34-1.61% @ SGI R4400 CPU; 20% @ SGI R4400 CPU	14 MB; 13.6 MB		22.8 KB startup, 1140 B typically; 1.68/7.36 & 60/2.76
	UNIX ISP Client	.56% @ SGI R4400 CPU	.55 MB		5.5/.253 & 5.2/.24
	Time Server				
	PRD Load Utility				
	Record Stream Data	15% @ 15 VUP[5]	64 MB	60 I/Os per sec	64 KB/sec
	Manage CCS DB (ORACLE)	5% @ 15 VUP	10 MB	20 I/Os per sec	
	Merge TLM Files	15% @ 15 VUP	64 MB	150 I/Os per sec	32 KB/sec
	Archive Stream Data	1% @ 15 VUP	2 MB	20 I/Os per sec	
	Manage Data Requests	1% @ 15 VUP	2 MB	10 I/Os per sec	1 KB/sec
	Perform Data Retrieval	15% @ 15 VUP per retrieval	8 MB per retrieval	200 I/Os per sec[6]	128 KB/sec
not Rel-1	Data (RedBrick) - Load	30 min x 4 CPUs	512 MB	1 GB	35 MB/day
not Rel-1	Data Warehouse (RedBrick) - Query	10 sec @ SGI R4400 CPU[7]	512 MB	35 MB/day	1 KB/request x 50 request/day +700 KB/response x 50 responses/day
not Rel-1	HSM (AMASS)	**not Rel-1**	**not Rel-1**	**not Rel-1**	**not Rel-1**
USER WS	Browser - Data Analysis	3% (25% for applet xfer) @ SGI R4400 CPU	30 MB	300 KB	300 KB
	Browser - RT display	50% @ SGI R4400 CPU	37 MB	200 KB	200 KB + .5 KB/sec

[4] Two sets of input: one for 3k & one for 6k mnemonics.
[5] VUP is VAX Unit of Processing; processing rate of a VAX 11/780 = 1 VUP.
[6] Disk retrieval data size range (512B, 50GB); typical size = 60 K Blocks x 512B.
[7] Estimate derived from RedBrick statistics for elapsed time & # physical I/Os for an example query against 1M rows of TLM data.

TEMPLATE 2: SOFTWARE RESOURCE REQUIREMENTS CHARACTERIZATION

(page 3 of 3)

Node	**Software CI**	**Processing Demand**	**Memory**	**Disk I/O Demand**	**Network Demand**
WEB SERVER	ISP Server[8]	1.34-1.61% @ SGI R4400 CPU; 8 % @ SGI R4400 CPU	14 MB; 8.2 MB	none	22.8 KB startup 1140 B typically; 4/.184
	UNIX ISP Client	.56 % @ SGI R4400 CPU	.55 MB		5.5./.254 & 5.2/.24
	JAVA ISP Client	.56 % @ SGI R4400 CPU	.55 MB		5.5./.254 & 5.2/.24
	Web Server- Data Analysis	1% @ SGI R4400 CPU	14 MB	300 KB	300 KB
	Web Server - RT display	1% @ SGI R4400 CPU	14 MB	200 KB	200 KB + .5 KB/sec
	GUI Server- Data Analysis	1% @ SGI R4400 CPU	12 MB	none	5 KB
	GUI Server- RT display	1% @ SGI R4400 CPU	12 MB	10 KB	10 KB
APPLICATION SERVER	Analysis Job Management	5% @ 24 VUP	4 MB	1 MB	none
	Build Analysis Request	5% @ 24 VUP	4 MB	1 MB	1 KB/sec
	Generate Analysis Product	12% @ 24 VUP	4 MB	none	none
	Manage Analysis Data	1% @ 24 VUP	1 MB	none	1 KB/sec
	Convert & Deliver Analysis Product	5% @ 24 VUP	2 MB	64 MB	64 MB/sec
	PVWAVE	3% @ 24 VUP	8 MB	64 MB	none

[8] Two sets of input: one for 3k & one for 6k mnemonics.

TEMPLATE 3: COTS OVERHEAD MAPPING TO RELEASE 1 HARDWARE

RELEASE 1 COMPUTERS

Overhead Function	FEP	Core Data Server	Core Firewall	Backbone Data Server	Web Server	User WS	Application Server
HP Openview	B	B	B	B	B	B	B
SNMP	B	B	B	B	B	B	B
ISIS/NET.H++	17, 17a, 18	17, 18		17, 17a, 18			
ISP	17a		17a	17a - Server	17a - Server	17a - Client	
NTP		B	B	B	B	B	B
IP	17, 17a, 18	17, 18	17, 17a, 18	7, 17, 17a, 18	7	7	7
TCP	17, 17a, 18	17, 18	17, 17a, 18	7, 17, 17a, 18	7	7	7
ORACLE		17, 18		7, 17, 18			
PV WAVE							7
JAVA						7	
FTP				7			7
IRIX		B		B	B	B	B
SOLARIS			B			B	
NETSCAPE					7 - Server	7 - Navigator	
HP Unix	B						
Link Layer	17, 17a, 18	17, 17a, 18	17, 17a, 18	7, 17, 17a, 18	7	7	7
Rogue Wave - DB Tools		17, 18		17, 18			
Rogue Wave - Tools	17, 17a, 18			17, 17a, 18			
VEDA	17, 18						
Dynamic Process Creation				7	17a		7

NOTES:

1. Thread 1 - Time.
2. Thread 7 - Data Analysis Request.
3. Thread 17 - Real Time Data Stream.
4. Thread 17a - Real Time Display (ISP).
5. Thread 18 - Recorded Data Stream.
6. B denotes a background overhead (i.e., not directly on the thread path).
7. RT WORKS, RED BRICK, & AMASS not in Release 1.
8. FEP Configuration Workstation, Network Management Node, & CM Server will not be modeled in the first iteration.

TEMPLATE 4: ORACLE CHARACTERIZATION FOR RELEASE 1

ORACLE Transaction Type	**Processing Rqmts per ORACLE Transaction**	**# Physical I/Os per ORACLE Transaction**	**Amount of Data per Physical I/O**
Telemetry Catalog Query	5 sec/request	15	40 B/ request
File Mgt Catalog Query	5 sec/ request	15	40 B/ request
PRD Data Query	5 sec/ request	781	512 B/ request

REDBRICK Transaction Type	**Processing Requirements per ORACLE Transaction**	**# Physical I/Os per ORACLE Transaction**	**Amount of Data per Physical I/O**
TLM Query	10 sec/trans	1367	512 B